

(11) **EP 0 856 797 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
21.05.2003 Bulletin 2003/21

(51) Int Cl.7: **G06F 12/08**

(21) Application number: **98300515.8**

(22) Date of filing: **26.01.1998**

(54) **A cache system for concurrent processes**

Cachespeichersystem für gleichzeitig laufende Prozesse

Système d'antémémoire pour des processus concurrents

(84) Designated Contracting States:
DE FR GB IT

(30) Priority: **30.01.1997 GB 9701960**
01.12.1997 GB 9725437

(43) Date of publication of application:
05.08.1998 Bulletin 1998/32

(73) Proprietor: **STMicroelectronics Limited**
Bristol BS12 4SQ (GB)

(72) Inventors:
• **Sturges, Andrew Craig**
Bath BA1 5NB (GB)
• **May, David**
Clifton, Bristol BS8 2EJ (GB)

(74) Representative: **Driver, Virginia Rozanne et al**
Page White & Farrer
54 Doughty Street
London WC1N 2LS (GB)

(56) References cited:
GB-A- 2 311 880

- **MUELLER F: "COMPILER SUPPORT FOR SOFTWARE-BASED CACHE PARTITIONING" ACM SIGPLAN NOTICES, vol. 30, no. 11, 1 November 1995, pages 125-133, XP000539037**
- **DONGWOOK K ET AL: "A PARTITIONED ON-CHIP VIRTUAL CACHE FOR FAST PROCESSORS" JOURNAL OF SYSTEMS ARCHITECTURE, vol. 43, no. 8, May 1997, pages 519-531, XP000685730**
- **KIRK D B: "SMART (STRATEGIC MEMORY ALLOCATION FOR REAL-TIME) CACHE DESIGN" PROCEEDINGS OF THE REAL TIME SYSTEMS SYMPOSIUM, SANTA MONICA, DEC. 5 - 7, 1989, no. SYMP. 10, 5 December 1989, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 229-237, XP000091005**

EP 0 856 797 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 0 856 797 B1

Description

[0001] The present invention relates to a cache system for operating between a processor and a main memory of a computer, and is particularly concerned with a processor capable of executing a plurality of concurrent processes.

[0002] As is well known in the art, cache memories are used in computer systems to decrease the access latency to certain data and code and to decrease the memory bandwidth used for that data and code. A cache memory can delay, aggregate and reorder memory accesses.

[0003] A cache memory operates between a processor and a main memory of a computer. Data and/or instructions which are required by the process running on the processor can be held in the cache while that process runs. An access to the cache is normally much quicker than an access to main memory. If the processor does not locate a required data item or instruction in the cache memory, it directly accesses main memory to retrieve it, and the requested data item or instruction is loaded into the cache. There are various known systems for using and refilling cache memories.

[0004] In order to rely on a cache in a real time system, the behaviour of the cache needs to be predictable. That is, there needs to be a reasonable degree of certainty that particular data items or instructions which are expected to be found in the cache will in fact be found there. Most existing refill mechanisms will normally always attempt to place in the cache a requested data item or instruction. In order to do this, they must delete other data items or instructions from the cache. This can result in items being deleted which were expected to be there for later use. This is particularly the case for a multi-tasking processor, or for a processor which has to handle interrupt processes or other unpredictable processes.

[0005] By way of background, reference is made to an article by Mueller, F. entitled "Compiler Support for Software-Based Cache Partitioning", ACM SIGPLAN NOTICES, vol. 30, no. 11, 1 November 1995, pages 125 to 133. This document discloses a scheme which partitions a cache according to tasks. The code and data of a task has to be restricted only to certain memory portions that map into the cache lines assigned to the task. Thus, if the code/data size of a task exceeds its cache partition size, the code/data has to be spread over the address space.

[0006] The scheme thus requires the compilers and linkers to organise code and data in a suitable manner.

[0007] GB 2311880 also describes a data processing system with a partitioned cache. In the particular case the cache is partitioned between instructions words and data words using a programmable partition setting held in a register.

[0008] Reference is also made to an article by Dongwook K. et al entitled "A portioned on-chip virtual cache for fast processors." Journal of System Architecture, vol. 43, no. 8, May 1997, pages 519 to 531. The article describes a cache which has a plurality of cache banks. Each process mainly accesses an assigned cache bank called a home bank. The identification of the home bank is stored in the process control block for each process. Although the process primarily accesses the home bank, in the system which is described in this article, a process can access other cache banks if needed. Thus, the system has a lower level of predictability than other more rigidly partitioned systems.

[0009] It is an object of the present invention to provide a cache system which provides greater predictability of caching behaviour for a processor executing a plurality of concurrent processes.

[0010] In this context, concurrent processes are considered to be processes which are executed by a common processor, but not necessarily simultaneously. That is, a first process may start to run and may be interrupted for some reason. The processor will then start to execute a second process but is ready to interrupt that when the first process is ready to run again or in response to some other prompt. This is managed by a process handler. It is important that data and/or instructions associated with the first process are not evicted from the cache while the second process is running. Conversely, it is useful to allow the second process to have access to the cache while it is running. Consider for example the situation illustrated in Figure 8 where two processes, process A and process B are running concurrently on one CPU. Process A is scheduled first and while it has the CPU it may completely fill the data cache with its own data, evicting any data which has been placed in the data cache for process B. When control then swaps to process B, it may then reverse the state of the data cache, throwing out all of the data of process A and bringing in its own. This ping-ponging of data cache state is common between concurrent processes and is often detrimental to performance.

[0011] According to one aspect of the present invention there is provided a method of operating a cache memory as set forth in claim 1.

[0012] By allocating a partition indicator to each process, processes running concurrently on the processor are prevented from evicting each others data and/or instructions from the cache memory. That is, the cache partition allocated for example to a first process running on the processor cannot be overwritten by a subsequent, second process. Instead, the second process will have its own cache partition allocated to it. It is of course preferable that the allocated partition indicators to processes can be altered so that once the first process has completely finished, the cache partition which was allocated to it can then be allocated to another process.

[0013] Depending on the needs of the process, it is possible to allocate more than one cache partition to a process or to deny a process access to the cache at all.

EP 0 856 797 B1

[0014] In the described embodiment, the partition indicator for a current process which is being executed is held in a process status store which also holds status information about the process. This is referred to herein as the thread status word register. When a new process is to be executed by the processor, a new thread status word is loaded into the store with a new partition indicator allocated to that process.

[0015] The partition indicator can be included in a group identifier for the process, the group identifier identifying an address space for the process. In a virtual addressing system, the processor issues addresses comprising a virtual page number and a line-in page number and a translation look-aside buffer is provided for translating the virtual page number to a real page number for accessing the main memory. The translation look-aside buffer can thus also receive the group identifier and derive therefrom the partition indicator for the current process depending on the virtual address space which has been allocated to the process.

[0016] The line-in page number of the items addressed can be used to identify the address storage location within the cache partition into which the item is to be located. That is, each cache partition is direct-mapped. It will be apparent that it is not necessary to use all of the end bits of the item's address as the line-in page number, but merely a set of appropriate bits. These will normally be near the least significant end of the address.

[0017] One or more cache partitions may be allocated to a process.

[0018] The system can include a cache access circuit which accesses items from the cache memory according to the address in main memory of said items and regardless of the cache partition in which the items is held in the cache memory. That is, the partition indicator is only used on refill and not on look-up. Thus, a cached item can be retrieved from its partition even if subsequent to its caching that partition is now allocated to a process associated with a different address space.

[0019] According to another aspect of the invention there is provided a computer systems as set forth in claim 8.

[0020] Each process can include one or more sequence of instructions held at addresses in the main memory within a common page number. Cache partitions can be allocated to processes by associating each cache partition with page numbers of a particular process in the main memory. This is described in our earlier GB Application No. 9701960.8.

[0021] As an alternative, a partition indicator can be held in the thread status word register and supplied directly to the cache refill mechanism.

[0022] The number of addressable storage locations in each cache partition can be alterable. Also, the association of cache partitions to page numbers can be alterable while a process using these page numbers is being run by the processor.

[0023] The following described embodiment illustrates a cache system which gives protection of the contents of the cache against unexpected eviction by reading from or writing to cache lines from other processes whose data are placed in other partitions. It also provides a system in which the contents of the cache may be predicted.

[0024] For a better understanding of the present invention and to show how the same may be carried into effect, reference will now be made by way of example to the accompanying drawings in which:

- Figure 1 is a block diagram of a computer incorporating a cache system;
- Figure 2 is a sketch illustrating a four way set associative cache;
- Figure 3 is a block diagram of the CPU of Figure 1;
- Figure 4 is an example of an entry in a translation look-aside buffer;
- Figure 5 is a block diagram of the refill engine;
- Figure 6 is a diagram illustrating the operation of a multi-tasking processor; and
- Figure 7 is a diagram illustrating the alteration in caching behaviour for the system of Figure 6;
- Figure 8 illustrates a non-partitioned cache; and
- Figures 9 and 10 illustrate useful applications of the invention.

[0025] Figure 1 is a block diagram of a computer incorporating a cache system. The computer comprises a CPU 2 which is connected to an address bus 4 for accessing items from a main memory 6 and to a data bus 8 for returning items to the CPU 2. Although the data bus 8 is referred to herein as a data bus, it will be appreciated that this is for the return of items from the main memory 6, whether or not they constitute actual data or instructions for execution by the CPU. The system described herein is suitable for use on both instruction and data caches. As is known, there may be separate data and instruction caches, or the data and instruction cache may be combined. In the computer described herein, the addressing scheme is a so-called virtual addressing scheme. The address is split into a line in page address 4a and a virtual page address 4b. The virtual page address 4b is supplied to a translation look-aside buffer (TLB) 10. The line in page address 4a is supplied to a look-up circuit 12. The translation look-aside buffer 10 supplies a real page address 14 converted from the virtual page address 4b to the look-up circuit 12. The look-up circuit 12 is connected via address and data buses 16,18 to a cache access circuit 20. Again, the data bus 18 can be for data items or instructions from the main memory 6. The cache access circuit 20 is connected to a cache memory 22 via an address bus 24, a data bus 26 and a control bus 28 which transfers replacement information for the cache memory 22. A refill

EP 0 856 797 B1

engine 30 is connected to the cache access circuit 20 via a refill bus 32 which transfers replacement information, data items (or instructions) and addresses between the refill engine and the cache access circuit. The refill engine 30 is itself connected to the main memory 6.

[0026] The refill engine 30 receives from the translation look-aside buffer 10 a full real address 34, comprising the real page address and line in page address of an item in the main memory 6. The refill engine 30 also receives a partition indicator from the translation look-aside buffer 10 on a four bit bus 36. The function of the partition indicator will be described hereinafter.

[0027] Finally, the refill engine 30 receives a miss signal on line 38 which is generated in the look-up circuit 12 in a manner which will be described more clearly hereinafter.

[0028] The cache memory 22 described herein is a direct mapped cache. That is, it has a plurality of addressable storage locations, each location constituting one row of the cache. Each row contains an item from main memory and the address in main memory of that item. Each row is addressable by a row address which is constituted by a number of bits representing the least significant bits of the address in main memory of the data items stored at that row. For example, for a cache memory having eight rows, each row address would be three bits long to uniquely identify those rows. For example, the second row in the cache has a row address 001 and thus could hold any data items from main memory having an address in the main memory which ends in the bits 001. Clearly, in the main memory, there would be many such addresses and thus potentially many data items to be held at that row in the cache memory. Of course, the cache memory can hold only one data item at that row at any one time.

[0029] Operation of the computer system illustrated in Figure 1 will now be described but as though the partition indicator was not present. The CPU 2 requests an item from main memory 6 using the address in main memory and transmits that address on address bus 4. The virtual page number is supplied to the translation look-aside buffer 10 which translates it into a real page number 14 according to a predetermined virtual to real page translation algorithm. The real page number 14 is supplied to the look-up circuit 12 together with the line in page number 4a of the original address transmitted by the CPU 2. The line in page address is used by the cache access circuit 20 to address the cache memory 22. The line in page address includes a set of least significant bits (not necessarily including the end bits) of the main address in memory which are equivalent to the row address in the cache memory 22. The contents of the cache memory 22 at the row address identified by the line in page address, being a data item (or instruction) and the address in main memory of the data item (or instruction), are supplied to the look-up circuit 12. There, the real page number of the address which has been retrieved from the cache memory is compared with the real page number which has been supplied from the translation look-aside buffer 10. If these addresses match, the look-up circuit indicates a hit which causes the data item which was held at that row of the cache memory to be returned to the CPU along data bus 8. If however the real page number of the address which was held at the addressed row in the cache memory 22 does not match the real page number supplied from the translation look-aside buffer 10, then a miss signal is generated on line 38 to the refill engine 30. It is the task of the refill engine 30 to retrieve the correct item from the main memory 6, using the real address which is supplied from the translation look-aside buffer 10 on bus 34. The data item, once fetched from main memory 6 is supplied to the cache access circuit 20 via the refill bus 32 and is loaded into the cache memory 22 together with the address in main memory. The data item itself is also returned to the CPU along data bus 8 so that the CPU can continue to execute. In a direct mapped cache memory as outlined above, it will be apparent that the data item and its address recalled from the main memory 6 will be loaded into the storage location from which the data item was originally accessed for checking. That is, it will be over-written into the only location which can accept it, having a row address matching the set of least significant bits in the line in page address in main memory. Of course, the page number of the data item originally stored in the cache memory and the data item which is now to be loaded into it are different. This "one to one mapping" limits the usefulness of the cache.

[0030] To provide a cache system with greater flexibility, an n-way set associative cache memory has been developed. An example of a 4-way set associative cache is illustrated in Figure 2. The cache memory is divided into four banks B1, B2, B3, B4. The banks can be commonly addressed row-wise by a common row address, as illustrated schematically for one row in Figure 2. However, that row contains four cache entries, one for each bank. The cache entry for bank B1 is output on bus 26a, the cache entry for bank B2 is output on bus 26b, and so on for banks B3 and B4. Thus, this allows four cache entries for one row address (or line in page address). Each time a row is addressed, four cache entries are output and the real page numbers of their addresses are compared with the real page number supplied from the translation look-aside buffer 10 to determine which entry is the correct one. If there is a cache miss upon an attempted access to the cache, the refill engine 30 retrieves the requested item from the main memory 6 and loads it into the correct row in one of the banks, in accordance with a refill algorithm which is based on, for example, how long a particular item has been held in the cache, or other program parameters of the system. Such replacement algorithms are known and are not described further herein.

[0031] Nevertheless, the n-way set associative cache (where n is the number of banks and is equal to four in Figure 2), while being an improvement on a single direct mapped system is still inflexible and, more importantly, does not allow the behaviour of the cache to be properly predictable.

EP 0 856 797 B1

[0032] The system described herein provides a cache partitioning mechanism which allows the optimisation of the computer's use of the cache memory by a more flexible cache refill system.

[0033] Figure 3 is a schematic block diagram of a CPU 2 using the computer of Figure 1. The CPU 2 comprises an execution circuit 15 which is connected to a fetch circuit 17 which is responsible for addressing memory via the memory bus 4 and retrieving data and instructions via the data bus 8. A set of general purpose registers 7 is connected to the execution circuit 15 for holding data and instructions for use in executing a process. In addition, a set of special registers are provided, denoted by reference numerals 9, 11 and 13. There may be any number of special purpose registers and by way of example register 11 holds the instruction pointer which identifies the line of code which is currently being executed. In addition, special register 9 holds a thread status word which defines the status of a process being executed by the CPU 2. The execution circuit 15 is capable of executing one process or sequence of instructions at any one time. However, it is equally capable of interrupting that process and starting to execute another process before the first process has finished executing. There are many reasons why a process currently being executed by the execution circuit 15 may be interrupted. One is that a higher priority interrupt process is to be executed immediately. Another is that the process being executed is currently awaiting data with a long latency, so that it is more efficient for the execution circuit to commence executing a subsequent process while the first process is waiting for that data. When the data has been received, the first process can be rescheduled for execution. The execution of concurrent processes is known *per se* and is managed by a process handler 19.

[0034] Each process is executed under a so-called "thread" of control. A thread has the following state:

- an instruction pointer which indicates where in the process the thread has advanced to,
- a jump pointer which indicates where the process will branch to next,
- a set of general purpose registers 7 which contain immediately accessible values,
- the mapping of virtual addresses to physical addresses,
- the contents of memory accessible through the virtual addresses,
- control registers accessible by the thread, and
- optionally other values such as floating point rounding mode, whether the thread has kernel privileges etc.

[0035] Some of the above state is specified by a small set of values which is referred to herein as thread status word and which is held in the thread status word register 9. The thread status word specifically holds information about:

- whether the thread is in kernel mode or not,
- which virtual address space the thread can access,
- the floating point flags, trap enables and modes,
- debug information, and
- trap optimisation information.

[0036] The format of the thread status word is defined in Table I.

TABLE I

Name	Bits ^a	Size	Description
TSW.FPFLAG	0-7	8	Floating point exception flags.
TSW.FPTRAP	8-15	8	Floating point exception traps.
TSW.FPMODE	16-19	4	Floating point modes.
	20-31		Reserved.
TSW.USER	32	1	Kernel mode (0)/user mode (1).

EP 0 856 797 B1

TABLE I (continued)

Name	Bits ^a	Size	Description
TSW.SINGLE	33	1	Single step mode.
TSW.TLB	34	1	First level TLB miss handler indicator.
TSW.WATCH	35	1	Watchpoints enabled.
TSW.ENABLE	36	1	Trap enable.
	37-47	11	Reserved.
TSW.GROUP	48-55	8	Group number.
	56-63		Reserved.

15 [0037] As can be seen from Table I, the thread status word includes an 8 bit group number. This is used as described in the following to generate the partition indicator for allocating cache partitions.

[0038] In the translation look-aside buffer 10 in the system described herein, each TLB entry has associated with the virtual page number, a real page number and an information sequence. The information sequence contains various information about the address in memory in a manner which is known and which will not be described further herein. However, according to the presently described system the information sequence additionally contains a partition code which generates a partition indicator PI dependent on the group number and the virtual page number. This is illustrated diagrammatically in Figure 4, where VP represents the virtual page number, RP represents the real page number, GN represents the group number and INFO represents the information sequence. In the described embodiment PI is four bits long.

25 [0039] Thus, bits 0 to 3 of the information sequence INFO constitute the partition indicator. The partition indicator gives information regarding the partition into which the data item may be placed when it is first loaded into the cache memory 22. For the cache structure illustrated in Figure 2, each partition can constitute one bank of the cache. In the partition indicator, each bit refers to one of the banks. The value of 1 in bit j of the partition indicator means that the data may not be placed in partition j. The value of 0 in bit j means that the data may be placed in partition j. Data may be placed in more than one partition by having a 0 in more than one bit of the partition indicator. A partition indicator which is all zeros allows the data to be placed in any partition of the cache. A partition indicator which is all ones does not allow any data items to be loaded into the cache memory. This could be used for example for "freezing" the contents of the cache, for example for diagnostic purposes.

30 [0040] In the example given in Figure 4, the partition indicator indicates that replacement of data items may not use banks B1 or B3 but may use banks B2 or B4.

35 [0041] It is quite possible to allocate more than one bank to a process. In that case, if the line in page address has more bits than the row address for the cache, the partitions would behave as a k-way set associative cache, where k partitions are allocated to a page. Thus, in the described example the process of Figure 4 can use banks B2 and B4. However, it may not use banks B1 and B3.

40 [0042] The partition information is not used on cache look-up, but only upon cache replacement or refill. Thus, the cache access can locate data items held anywhere in the cache memory, whereas a replacement will only replace data into the allowed partitions for that process.

45 [0043] Figure 5 illustrates in more detail the content of the refill engine 30. The refill bus 32 is shown in Figure 4 as three separate buses, a data bus 32a, an address bus 32b and a bus 32c carrying replacement information. The address and data buses 32a and 32c are supplied to a memory access circuit 50 which accesses the main memory via the memory bus 54. The replacement information is fed to a decision circuit 52 which also receives the real address 34, the partition indicator P1 on bus 36 and the miss signal 38. The decision circuit 52 determines the proper partition of the cache into which data accessed the main memory is to be located.

50 [0044] The cache partitioning mechanism described herein is particularly useful for a multi-tasking CPU. A multi-tasking processor is capable of executing concurrent processes, that is running more than one process "simultaneously". In practice, the processor executes part of a process and, when that process is halted for some reason, perhaps in need of data or a stimulus to proceed, the processor immediately begins executing another process. Thus, the processor is always operating even when individual processes may be held up waiting for data or another stimulus to proceed. Figure 6 illustrates diagrammatically such a situation. On the left hand side of Figure 6 is illustrated the sequence which a processor may undertake to run different processes P1, P2, P3, P4. On the right hand side of Figure 5 is an illustration of where these processes may expect their data to be held in memory. Thus, the data for the process P1 are held on page 0. The data for process P2 are held on pages 1 and 2. Data for processes P3 and P4 share page 3. In the example, the processor executes a first sequence of process P1, a first sequence of process P2, a second

EP 0 856 797 B1

sequence of process P1, a second sequence of process P2 and then a first sequence of process P3. When the second sequence of the process P1 has been executed, the process P1 has been fully run by the processor. It will readily be apparent that in a conventional cache system, once the processor has started executing the first sequence of the process P2, and is thus requesting accesses from page 1, the data items and instructions in these lines will replace in the cache the previously stored data items and instructions from page 0.

[0045] However, these may soon again be required when the second sequence of the process P1 is executed.

[0046] The cache partitioning mechanism described herein avoids the timing delays and uncertainties which can result from this. Figure 7 shows the partitioning of the cache while the processor is running process P1, and the change in the partitioning when the processor switches to running P3 etc. Figure 6 also shows the TLB cache partition indicators for each case. Thus, on the left hand side Figure 5 shows the cache partitioned while the processor is running processes P1 and P2. The process P1 may use banks B1 and B2 of the cache, but may not use banks B3 and B4. Conversely, the process P2 may use banks B3 and B4, but not banks B1 and B2. This can be seen in the TLB entries. This is, process P1 has a cache partition indicator allowing it to access banks B1 and B2, but not B3 and B4. Process P2 has cache partition indicators allowing them to access banks B3 and B4 but not B1 and B2. Process P3 has a cache partition indicator which prevents it from accessing the cache. Thus, any attempt by the processor to load data items from the process P3 into the cache would be prohibited. For the described process sequence, this however is not a disadvantage because, as can be seen, the processor is not intending to execute any part of the process P3 until it has finished executing process P1. If it did for some reason have to execute P3, the only downside would be that it would have to make its accesses from direct memory and would not be allowed use of the cache.

[0047] When the process P1 has finished executing, the processor can request kernel mode to allow it to alter the cache partition indicators for the processes. The manner in which this is done depends on how the partitioning mechanism is implemented. For the above described example, the partition code can be set in the TLB like any other TLB entry. Thus, the partition codes are normally set by kernel mode software running on the CPU 2. However, a user may alter partitions by requesting that the cache partitions be altered. In that event, the CPU 2 would change to kernel mode to implement the request, change the TLB entries accordingly and then return to the user mode to allow the user to continue. Thus, a user can alter the partitioning behaviour of the cache, thus providing much greater flexibility than has hitherto been possible. The change is illustrated on the right hand side of Figure 6. Thus, now the cache partition indicators prevent the process P1 from using the cache at all, but allocate banks B1 and B2 to the processes P3 and P4, by altering the cache partition indicator for processes P3 and P4 so that it can access these banks of the cache. Thus, when the processor is expecting to execute the process P3, it now has a cache facility.

[0048] Thus, a system has been described above which prevents concurrent processes from evicting each others data from the data cache. That is, the processes are mapped to disjoint data cache partitions. This effectively gives each process its own private data cache. While this does reduce the amount of data cache space available to both processes, it makes their performance much easier to predict accurately. The result of the system described herein is illustrated in Figure 9.

[0049] Another area where the system described herein is particularly useful is in the implementation of performance-critical routines. Often, there are a few routines whose performance is absolutely critical to the overall performance of the system. A good example of this might be an interrupt service routine which, when called, must produce an effect in a guaranteed (and usually short) length of time. In these cases, cache partitions may be reserved in both the data and instruction caches for the data and code required for these important routines. The rest of the instruction and data caches may then be shared out among the remaining processes. Figure 10 illustrates a possible arrangement. In Figure 10, an example is shown of reserving one 2 kilobyte partition in the data cache and one 4 kilobyte partition in the instruction cache for a performance critical interrupt service routine.

[0050] It will be appreciated that the present invention is not restricted to the specifically described embodiment above. Some particular possible variations are mentioned below, but this is not a comprehensive list of the variations which are possible within the scope of the invention.

[0051] In the described embodiment above, the address issued by the CPU on address bus 4 is split into a virtual page number 4b and a line in page 4a. However, the invention can also be used in a situation where the entire virtual address is sent from the CPU to the look-up circuit for the cache. Conversely, the invention is also applicable in a situation where the CPU issues real addresses directly to the look-up circuit. What is important is that the cache partition indicator is provided in association with the particular process being executed.

[0052] In the embodiment described above, a single cache access circuit 20 is shown for accessing the cache both on look-up and refill. However, it is also possible to provide the cache with an additional access port for refill, so that look-up and refill take place via different access ports for the cache memory 22.

[0053] In the described embodiment, the refill engine 30 and cache access circuit 20 are shown in individual blocks. However, it would be quite possible to combine their functions into a single cache access circuit which performs both look-up and refill.

[0054] The following are possible alternatives for generation of the partition indicator in association with a particular

EP 0 856 797 B1

process.

[0055] In one alternative, the partition indicator may be placed directly in the thread status word TSW. For the described thread status word, this could be done by allocating the currently reserved bits 56 to 59 of the TSW to a new field TSW.PI. The value of the TSW.PI would then be passed to the refill engine 30 directly from the CPU. This would require a modification to the architecture illustrated in Figure 1 to connect the partition indicator bus 36 directly from the CPU 2 to the refill engine 30, rather than from the TLB 10. For this implementation, the partition indicator PI is changed when a new thread status word TSW is loaded for the next thread to be executed. This can be done by a particular setting instruction which sets the parameters of the thread status word.

[0056] In another implementation, a table can be provided in the translation look-aside buffer 10 which maps the group number to the partition indicator PI without using the virtual address. This could be done by having a table indexed by group number which returns the partition indicator, or a table which has a group number/partition indicator pair which returns the partition indicator for the matching group. In this case, the architecture of Figure 1 would be unaltered, but a different table would be required in the translation look-aside buffer 10. In this implementation, the partition indicator could be changed using a "put" instruction which has two operands, the control register number to be changed and its new value. All control registers would be allocated a number which can be used to access it using this instruction, and so each entry in the group number/partition indicator table would have a unique control register number.

Claims

1. A method of operating a cache memory (22) arranged between a processor (2) and a main memory (6) of a computer, the processor being capable of executing a plurality of processes wherein each process includes a sequence of instructions, the method comprising:

dividing the cache memory into cache partitions (B1,B2,B3,B4), each cache partition having a plurality of addressable storage locations for holding items in the cache memory; and

when the processor requests an item from main memory during execution of a current process and that item is not held in the cache memory, fetching the item from main memory and loading it into one of the plurality of addressable storage locations in a cache partition identified for use by that process;

characterised by allocating to each process a partition indicator (PI) identifying which, if any, of said cache partitions is to be used for holding items for-use in the execution of that process, and wherein:

the partition indicator comprises a plurality of bits associated respectively with the partitions, the status of each bit indicating whether or not that process can use the partition associated with that bit;
upon a cache miss, the item is loaded on refill only into a cache partition identified by the partition indicator as useable by that process; and
items are accessed on lookup from the cache memory regardless of the cache partition in which the item is held in the cache memory.

2. A method according to claim 1, comprising the step of holding in a store (9) the partition indicator for the current process which is currently being executed.

3. A method according to claim 2, wherein when a new process is to be executed by the processor, a new partition indicator allocated to that new process is loaded into said store.

4. A method according to claim 2 or 3, wherein the store (9) which holds the partition indicator for the current process is a process status store which also holds status information about the process.

5. A method according to claim 1, 2 or 3, wherein the partition indicator is included in a group identifier for the process, the group identifier identifying an address space for the process.

6. A method according to claim 5, wherein the processor issues addresses comprising a virtual page number (VP) and a line-in page number and wherein a translation look-aside buffer (10) is provided for translating the virtual page number to a real page number (RP) for accessing the main memory, the translation look-aside buffer also receiving the group identifier and deriving therefrom the partition indicator for the current process.

EP 0 856 797 B1

7. A method according to any preceding claim, wherein the number of addressable storage locations in each cache partition is alterable.

8. A computer system comprising:

5

a processor (2) for executing a plurality of processes wherein each process includes a sequence of instructions;
a main memory (6);

a cache memory (22) having a set of cache partitions, each cache partition comprising a plurality of addressable storage locations for holding items fetched from said memory for use by the processor in execution of its processes; and

10

a cache refill mechanism (30) arranged to fetch an item from the main memory and to load said item into the cache memory at one of said addressable storage locations, **characterised in that** the processor includes a process status store which holds a partition indicator (PI) for a current process which is currently being executed, the partition indicator comprising a plurality of bits associated respectively with the partitions, the status of each bit indicating whether or not that process can use the partition associated with that bit;

15

the cache refill mechanism selects upon a cache miss, said one of said addressable storage locations for loading said items only into the cache partition(s) identified by the partition indicator held in the process status store in association with the current process; and by

20

a cache access circuit (20) operable on lockup to access items from the cache memory regardless of the cache partition in which the item is held in the cache memory.

9. A computer system according to claim 8, wherein the partition indicator (PI) is included in a group identifier (GN) for each process which identifies an address space for the process.

25

10. A computer system according to claim 9, wherein the processor issues addresses comprising a virtual page number (VP) and a line-in page number and wherein the system comprises a translation look-aside buffer (10) for translating the virtual page number to a real page number (RP) for accessing the main memory, the translation look-aside buffer being operable to receive said group identifier and to derive therefrom the partition indicator for the current process.

30

Patentansprüche

1. Verfahren zum Betreiben eines Cachespeichers (22), der zwischen einem Prozessor (2) und einem Hauptspeicher (6) eines Rechners angeordnet ist, wobei der Prozessor in der Lage ist, eine Mehrzahl von Prozessen auszuführen, wobei jeder Prozeß eine Folge von Befehlen umfasst, wobei das Verfahren umfasst:

35

Aufteilen des Cachespeichers in Cachespeicherbereiche (B1, B2, B3, B4), wobei jeder Cachespeicherbereich eine Mehrzahl von adressierbaren Speicherplätzen zum Speichern von Daten in dem Cachespeicher aufweist; und

40

wenn der Prozessor ein Datenelement vom Hauptspeicher während der Ausführung eines gegenwärtigen Prozesses anfordert und dieses Datenelement nicht im Cachespeicher enthalten ist, Abrufen des Datenelements vom Hauptspeicher und Laden desselben in einen der Mehrzahl der adressierbaren Speicherplätze in einem Cachespeicherbereich, der für die Verwendung durch diesen Prozeß identifiziert ist;

45

gekennzeichnet durch Zuweisen eines Speicherbereichsindikators (PI) an jeden Prozeß, welcher anzeigt, welcher der Cachespeicherbereiche zum Speichern von Daten zur Verwendung bei der Ausführung dieses Prozesses anzuwenden ist, und wobei:

50

der Speicherbereichsindikator eine Mehrzahl von Bits aufweist, die jeweils den Speicherbereichen zugeordnet sind, wobei der Status von jedem Bit anzeigt, ob dieser Prozeß den diesem Bit zugeordneten Speicherbereich verwenden kann;

bei einem Cachefehlgriß das Datenelement beim Nachfüllen nur in einen Cachespeicherbereich geladen wird, welcher **durch** den Speicherbereichsindikator als für diesen Prozeß verwendbar gekennzeichnet ist; und

55

auf Daten vom Cachespeicher in einem Lookup-Vorgang zugegriffen wird, unabhängig von dem Cachespeicherbereich, in welchem das Datenelement im Cachespeicher gespeichert ist.

2. Verfahren nach Anspruch 1, bei welchem in einem Speicher (9) der Speicherbereichsindikator für den gegenwär-

EP 0 856 797 B1

tigen Prozeß, der gegenwärtig ausgeführt, gespeichert wird.

3. Verfahren nach Anspruch 2, bei welchem, wenn ein neuer Prozeß durch den Prozessor ausgeführt werden soll, ein neuer Speicherbereichsindikator, welcher diesem neuen Prozeß zugewiesen ist, in den Speicher geladen wird.

4. Verfahren nach Anspruch 2 oder 3, bei welchem der Speicher (9), welcher den Speicherbereichsindikator für den gegenwärtigen Prozeß beinhaltet, ein Prozeßstatusspeicher ist, der auch eine Statusinformation über den Prozeß speichert.

5. Verfahren nach Anspruch 1, 2 oder 3, bei welchem der Speicherbereichsindikator in einem Gruppenkennzeichen für den Prozeß enthalten ist, wobei das Gruppenkennzeichen einen Adressbereich für den Prozeß identifiziert.

6. Verfahren nach Anspruch 5, bei welchem der Prozessor Adressen erzeugt, welche eine virtuelle Seitenzahl (VP) und eine Zeilenseitenzahl umfassen, und wobei ein paralleler Umsetzungspuffer (10) zum Umsetzen der virtuellen Seitenzahl in eine reelle Seitenzahl (RP) zum Zugreifen auf den Hauptspeicher vorgesehen ist, wobei der Umsetzungspuffer auch das Gruppenkennzeichen empfängt und daraus den Speicherbereichsindikator für den gegenwärtigen Prozeß ableitet.

7. Verfahren nach einem der vorstehenden Ansprüche, bei welchem die Anzahl der adressierbaren Speicherplätze in jedem Cachespeicherbereich veränderlich ist.

8. Rechnersystem, aufweisend:

einen Prozessor (2) zum Ausführen von einer Mehrzahl von Prozessen, wobei jeder Prozeß eine Folge von Befehlen umfasst;

einen Hauptspeicher (6) ;

einen Cachespeicher (22) mit einer Gruppe von Cachespeicherbereichen, wobei jeder Cachespeicherbereich eine Mehrzahl von adressierbaren Speicherplätzen zum Speichern von Daten aufweist, die von dem Speicher zur Verwendung durch den Prozessor bei der Ausführung seiner Prozesse abgerufen werden;

einen Cachenachfüllmechanismus (30), der angeordnet ist, um ein Datenelement aus dem Hauptspeicher abzurufen und dieses Datenelement in den Cachespeicher in einen der adressierbaren Speicherplätze zu laden, **dadurch gekennzeichnet, dass** der Prozessor einen Prozeßstatusspeicher aufweist, welcher einen Speicherbereichsindikator (PI) für einen gegenwärtigen Prozeß speichert, der gegenwärtig ausgeführt wird, wobei der Speicherbereichsindikator eine Mehrzahl von Bits aufweist, die jeweils den Speicherbereichen zugeordnet sind, wobei der Status von jedem Bit anzeigt, ob dieser Prozeß den diesem Bit zugeordneten Speicherbereich verwenden kann;

der Cachenachfüllmechanismus bei einem Cachefehlgriß den einen der adressierbaren Speicherplätze zum Laden der Daten nur in den/die Cachespeicherbereich(e) auswählt, der/die durch den Speicherbereichsindikator, der in Verbindung mit dem gegenwärtigen Prozeß im Prozeßstatusspeicher gespeichert ist, angegeben wird/werden; und durch

eine Cachezugriffsschaltung (20), die in einem Lookup-Vorgang betreibbar ist, um auf Daten aus dem Cachespeicher zuzugreifen, unabhängig von dem Cachespeicherbereich, in welchem das Datenelement im Cachespeicher gespeichert ist.

9. Rechnersystem nach Anspruch 8, bei welchem der Speicherbereichsindikator (PI) in einem Gruppenkennzeichen (GN) für jeden Prozeß enthalten ist, welches einen Adressbereich für den Prozeß angibt.

10. Rechnersystem nach Anspruch 9, bei welchem der Prozessor Adressen erzeugt, welche eine virtuelle Seitenzahl (VP) und eine Zeilenseitenzahl umfassen, und wobei ein paralleler Umsetzungspuffer (10) zum Umsetzen der virtuellen Seitenzahl in eine reelle Seitenzahl (RP) zum Zugreifen auf den Hauptspeicher vorgesehen ist, wobei der Umsetzungspuffer betreibbar ist, um das Gruppenkennzeichen zu empfangen und daraus den Speicherbereichsindikator für den gegenwärtigen Prozeß abzuleiten.

Revendications

1. Procédé d'actionnement d'une mémoire cache (22) disposée entre un processeur (2) et une mémoire principale (6) d'un ordinateur, le processeur pouvant exécuter une pluralité de processus, chacune processus comprenant

EP 0 856 797 B1

une séquence d'instructions, le procédé comprenant les étapes suivantes :

diviser la mémoire cache en partitions de cache (B1, B2, B3, B4), chaque partition de cache comportant une pluralité d'emplacements de mémorisation adressables pour stocker des éléments dans la mémoire cache ; et quand le processeur requiert un élément de la mémoire principale pendant l'exécution d'un processus courant et que cet élément n'est pas stocké dans la mémoire cache, rechercher l'élément à partir de la mémoire principale et le charger dans l'un de la pluralité d'emplacements de mémorisation adressables dans une partition de cache identifiée pour être utilisée par ce processus ;

caractérisé par l'allocation à chaque processus d'un indicateur de partition (PI) identifiant laquelle, s'il y en a, des partitions de cache doit être utilisée pour stocker des éléments destinés à être utilisées dans l'exécution de ce processus, et dans lequel :

l'indicateur de partition comprend une pluralité de bits associés respectivement aux partitions, l'état de ce bit indiquant si le processus peut utiliser ou non la partition associée à ce bit ;
à la suite d'une absence dans le cache, l'élément est chargé lors d'un re-remplissage seulement dans la partition de cache identifiée par l'indicateur de partition comme utilisable pour ce processus ; et on accède aux éléments lors d'une consultation à partir de la mémoire cache quelle que soit la partition de cache dans laquelle l'élément est stocké dans la mémoire cache.

2. Procédé selon la revendication 1, comprenant l'étape de chargement dans la mémoire (9) de l'indicateur de partition pour le processus courant en cours d'exécution.

3. Procédé selon la revendication 2, dans lequel, quand un nouveau processus doit être exécuté par le processeur, un indicateur de partition alloué à ce nouveau processus est chargé dans la mémoire.

4. Procédé selon la revendication 2 ou 3, dans lequel la mémoire (9) qui maintient un indicateur de partition pour le processus courant est une mémoire d'état de processus qui stocke également des informations d'état sur le processus.

5. Procédé selon la revendication 1, 2 ou 3, dans lequel l'indicateur de partition est inclus dans un indicateur de groupe pour le processus, l'identificateur de groupe identifiant un espace d'adresse pour le processus.

6. Procédé selon la revendication 5, dans lequel le processeur fournit des adresses comprenant un numéro de page virtuel (VP) et un numéro de page en cours et dans lequel un tampon de consultation de traduction (10) est prévu pour traduire le numéro de page virtuel en un numéro de page réel (RP) pour accéder à la mémoire principale, le tampon de consultation de traduction recevant également l'identificateur de groupe et fournissant à partir de là l'indicateur de partition pour le processus courant.

7. Procédé selon l'une quelconque des revendications précédentes, dans lequel le numéro des emplacements de mémorisation adressable dans chaque partition de cache est modifiable.

8. Système d'ordinateur comprenant :

un processeur (2) pour exécuter une pluralité de processus, chaque processus comprenant une séquence d'instructions ;

une mémoire principale (6) ;

une mémoire cache (22) comportant un ensemble de partitions de cache, chaque partition de cache comprenant une pluralité d'emplacements de mémorisation adressables pour stocker des éléments recherchés dans la mémoire pour être utilisés par le processeur lors de l'exécution de ses processus ; et

un mécanisme de re-remplissage de cache (30) agencé pour rechercher un élément à partir de la mémoire principale et charger cet élément dans la mémoire cache à l'un des emplacements de mémorisation adressables,

caractérisé en ce que le processeur comprend une mémoire d'état de processus qui stocke un indicateur de partition (PI) pour un processus courant qui est en cours d'exécution, l'indicateur de partition comprenant une pluralité de bits associés respectivement aux partitions, l'état de chaque bit indiquant si ce processus peut utiliser ou non la partition associée à ce bit ;

EP 0 856 797 B1

le mécanisme de re-remplissage de cache sélectionne, à la suite d'une absence dans le cache, ledit un des emplacements de mémorisation adressables pour charger les éléments seulement dans la ou les partitions de cache identifiées par l'indicateur de partition stocké dans la mémoire d'état de processus en association avec le processus courant ; et par

5 un circuit d'accès en cache (20) actionnable lors d'une recherche pour accéder lors d'une consultation à des éléments en provenance de la mémoire cache quelle que soit la partition de cache dans laquelle l'élément est stocké dans la mémoire cache.

10 9. Système d'ordinateur selon la revendication 8, dans lequel l'indicateur de partition (PI) est inclus dans un indicateur de groupe (GN) pour chaque processus qui identifie un espace d'adresse pour le processus.

15 10. Système d'ordinateur selon la revendication 9, dans lequel le processeur fournit des adresses comprenant un numéro de page virtuel (VP) et un numéro de page en cours et dans lequel un tampon de consultation de traduction (10) est prévu pour traduire le numéro de page virtuel en un numéro de page réel (RP) pour accéder à la mémoire principale, le tampon de consultation de traduction recevant également l'identificateur de groupe et fournissant à partir de là l'indicateur de partition pour le processus courant.

20

25

30

35

40

45

50

55

EP 0 856 797 B1




FIG. 1

EP 0 856 797 B1

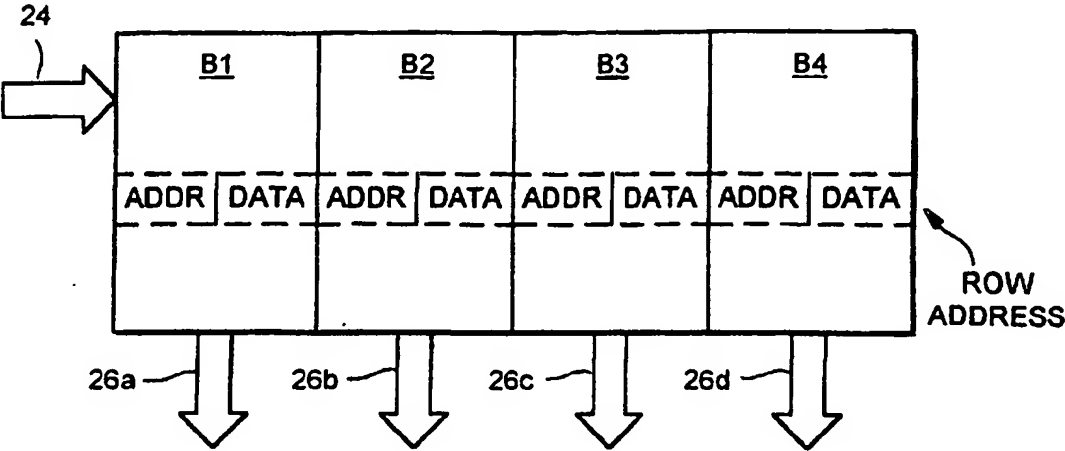


FIG. 2

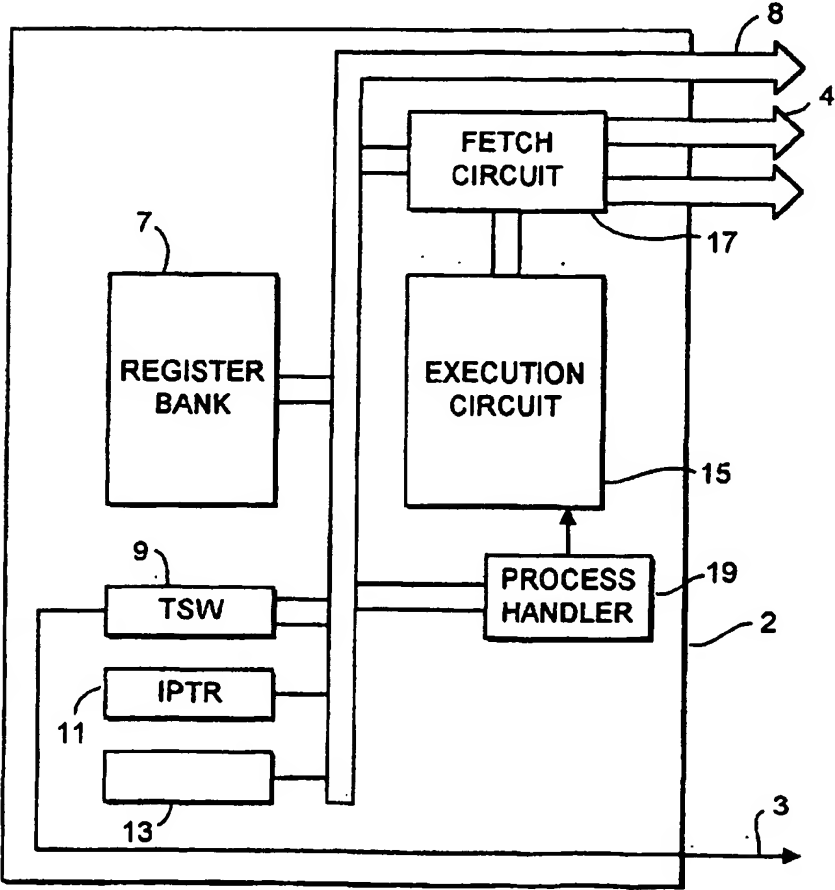
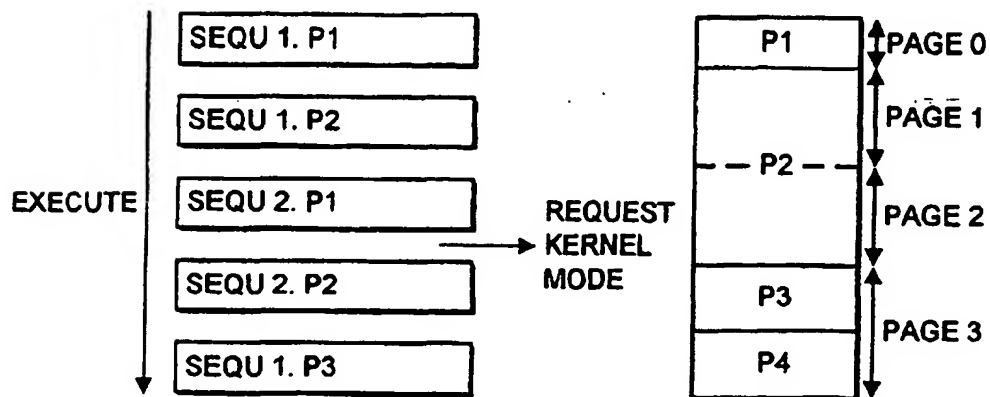

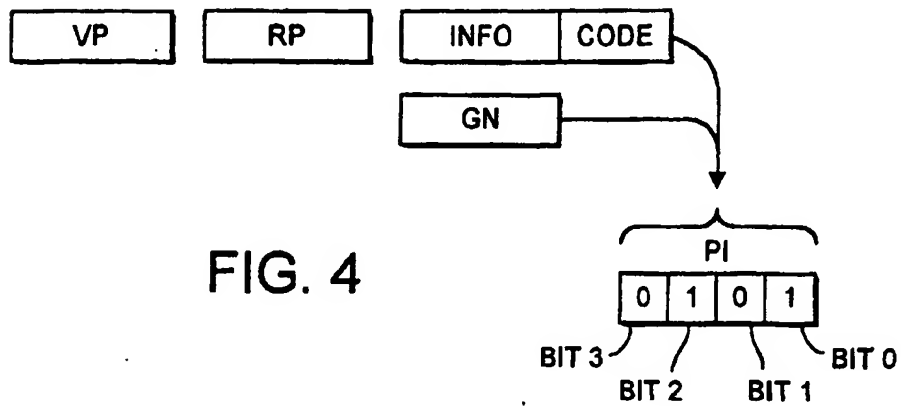


FIG. 3

EP 0 856 797 B1



EP 0 856 797 B1

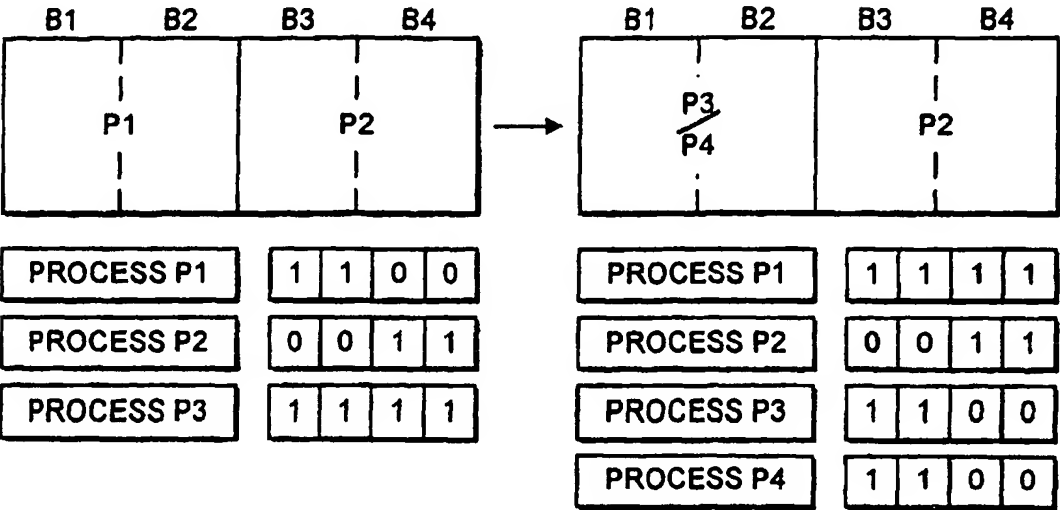


FIG. 7




FIG. 8

EP 0 856 797 B1




FIG. 9

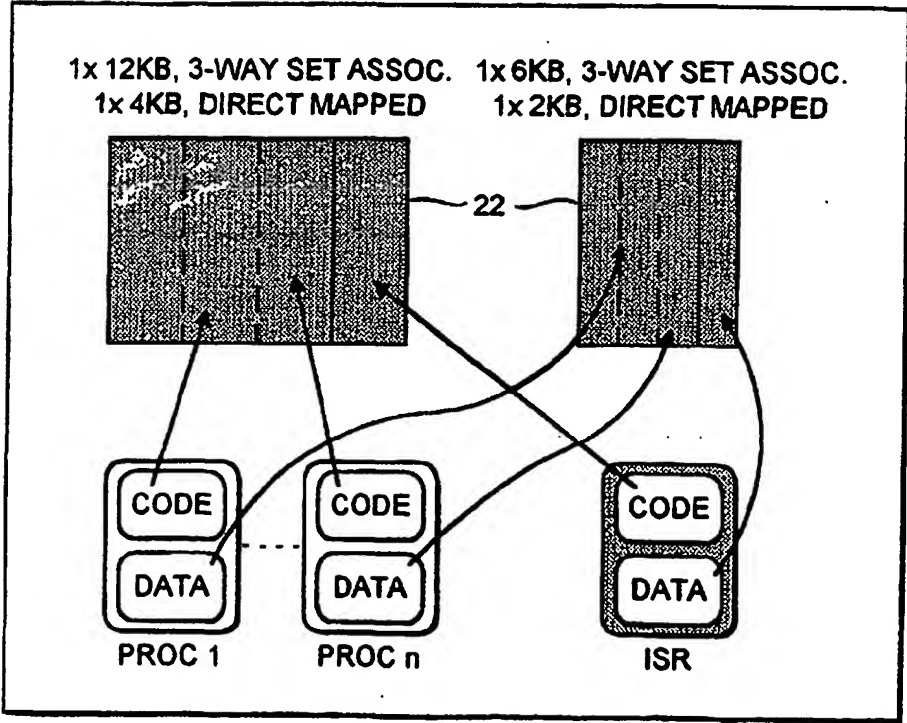


FIG. 10